



# Automotive cross-domain security framework

Ensuring safety, security, and flexibility  
in next-gen vehicle computers

# Abstract

In this white paper, we propose an “automotive cross-domain security” framework that involves defining distinct domains with strictly separated communication channels. Our approach ensures that functionalities remain isolated and yet interconnected, akin to different sites in an enterprise network, allowing for a secure and flexible architecture needed for next-gen vehicle apps. We give detailed insights into the structure of such an approach for the journey towards the software-defined vehicle, highlighting the mitigation of risks from unauthorized access. In addition to the technical description, we use a real-life attack scenario on a vehicle computer to demonstrate the actual benefits of combining trusted execution environments and hardware security modules. The paper further explains how the “automotive cross-domain security” framework can be used to establish a Zero Trust paradigm for the fully secure vehicle computer of the future.



# 1. Introduction

The vision of the automotive industry is set. Soon, fleets will consist of software-defined vehicles (SDV) that evolve permanently during their entire lifecycle through updates of apps and services. For manufacturers, this means short innovation cycles with high pressure to release new functionalities in a very competitive market while fulfilling all (cyber)security requirements. The balancing act between security and speed of innovation becomes apparent in the main component of the SDV: the vehicle computer. The goal is a system where dynamic adaptation without much integrational or programming effort and safety go hand in hand. An effective way to achieve this is through separated, yet interconnected domains within the vehicle computer, which allow the combination of safety-critical with non-critical functionalities without endangering vehicle safety.

Dealing with a growing number of external cybersecurity threats is not unique to the automotive industry. Wherever digitalization picks up speed, the issue of establishing high flexibility and connectivity within a complex system while maintaining functional integrity is an important topic. So, it makes sense to look at the bigger picture and to consider tried and tested solutions from other areas. A deep dive into enterprise IT architectures, for example, reveals the already established security principle of cross-domain solutions<sup>1</sup>. By transferring the lessons learned in this domain to the vehicle computer, we developed the "automotive cross-domain security" framework. The security aspect is covered by strictly separated communication channels and by the isolation of safety-relevant functionalities.

## 2. Challenge of vehicle computer security


In classical vehicle software development, the engineering process is relatively independent of the scope and depth of the application. In any case, tests are necessary to check the integrity of all vehicle functions, i.e. to ensure that even an apparently harmless app does not trigger malfunctions which could jeopardize driver safety or act as an entry point for malicious software. With the increasing number of access channels and requests from outside, the vehicle is exposed to more and more risks, and the effort required to maintain general vehicle security is rising exponentially.

This calls for new software architectures with simplified software development, deployment, and operation, which at the same time maintain a high level of user safety and system reliability. The vehicle computer as a centralized control point for various vehicle functionalities heralds the path to a more software-defined future. However, it also comes with challenges. Automotive manufacturers must achieve the right level of performance while ensuring appropriate security measures and high flexibility against a growing landscape of cyber threats. Many players in the automotive industry are breaking new ground here. It is worthwhile to think outside the box and adopt successful approaches instead of reinventing the wheel.



### Deep dive into automotive cybersecurity

Would you like to find out more about cybersecurity in the automotive industry? Our [white paper "Automotive cyber-security fully revealed"](#) is a holistic guide to navigating the highly dynamic world of automotive cybersecurity, diving deeper into the primary challenges and opportunities vehicle manufacturers are facing.



### 3. Safety-related functionalities

When thinking of safety-related functions in a vehicle, the first things that come to mind are the braking function, engine management, airbags, etc. Despite the changes in vehicle architectures on the way to the SDV, these highly safety-critical components often continue to be controlled separately, via separate ECUs and not as part of a vehicle computer domain. However, a whole range of functionalities and signals that directly affect these safety components will

inevitably be shifted to the vehicle computer. This starts with the control and processing of sensor technology and ends with sophisticated functions that influence driving, such as obstacle detection (or ADAS in general) or electronic stability control (ESC). When talking about safety-relevant functions in this paper, we refer to this large, partly newly emerging category, hosted on the vehicle computer together with entertainment, comfort, and connectivity functionalities.

### 4. Learning from enterprise IT security

The good news is that the automotive industry is not the first one to face the challenge of balancing flexibility with security within a complex and subsequently more connected system. Critical industries like energy, finance, and healthcare also have a growing IT landscape with various security levels. For example, an energy company might have internal networks that control grid operations (sensitive data) and external-facing networks for customer services (less sensitive data). Their cross-domain approach consists of using the principles of virtualization to split complex systems into isolated units with a security domain for control functionalities and real-time threat detection and prevention, as summarized in figure 1.

The cross-domain approach allows them to manage secure access and data transfer between different domains. It

makes sure that operational data remains protected while allowing external communication, for example through internet connectivity with bi-directional data flow to various sources. These industry players have already gone down this path and have found ways to counter the growing cyber threat landscape.

A single vehicle must be considered as a complex IT network corresponding to the network of a company. It has numerous domains that are or can be connected to each other internally, as well as with external sources. The principles mentioned above can thus be mapped to the vehicle's software – without losing sight of the fact that, despite all digitalization, it will always be a mechanical device intended to transport people. Hence, the well-being of all road users must always have top priority.

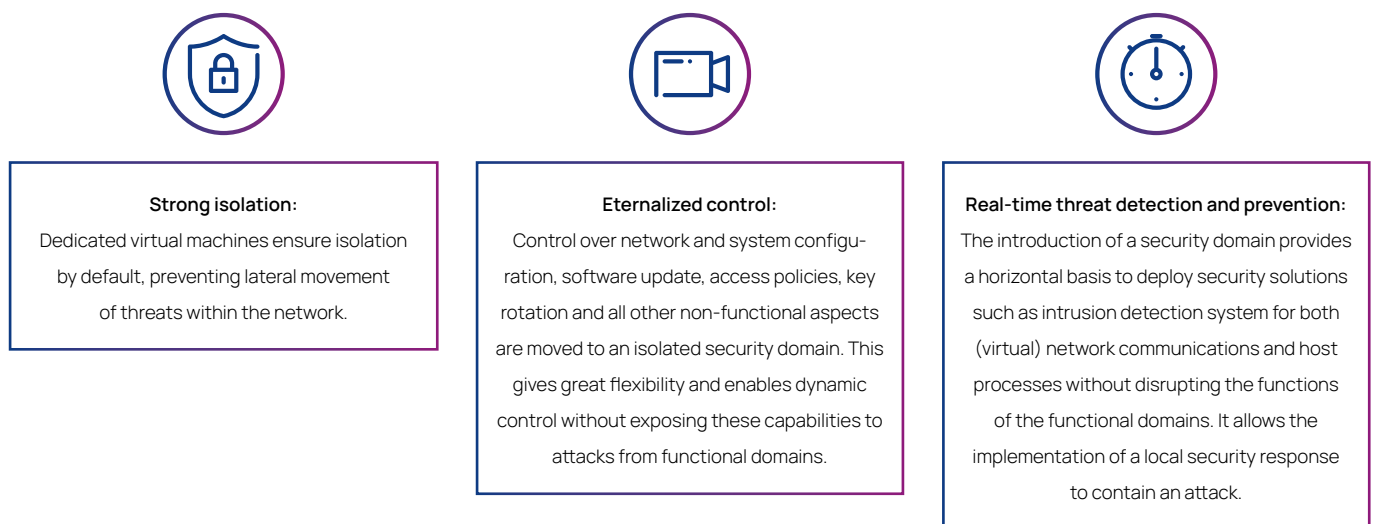


Figure 1: Principles within the IT industry for facing external cybersecurity threats and simplifying their IT by adopting a cross-domain approach.

## 5. Adopting learnings for the vehicle computer

To adopt the cross-domain approach within the automotive sector, we transfer both the basic pattern and the concept of using additional domains for management and connectivity purposes. The result is an “automotive cross-domain security” framework for the creation of differentiated and localized security responses by utilizing virtual network channels between virtual machines (VM)<sup>2</sup> and applications. They allow for the seamless connection of diverse domains – and facilitate flexible software architecture within the SDV’s software architecture.

Risks can be mitigated effectively by segregating domains and tailoring prevention strategies to their specific requirements. Our vehicle-specific approach presented in figure 2 consists of four domains.




Figure 2: The domains (virtual machine categories) of the automotive cross-domain security framework within a vehicle computer.

The **security domain** serves as a gateway for all communication. It incorporates robust security measures, including dedicated VMs, hardware security modules, and trusted execution environments. Moreover, communication within this domain undergoes rigorous analysis through firewall systems and intrusion detection mechanisms. The security domain holds exclusive authority over system configuration and policies, leveraging technologies like trusted execution environments to ensure separation from regular operations. This setup empowers the vehicle computer to implement proactive prevention mechanisms, thereby reducing reliance on external security operations centers and containing an attack until a fix is available.

The **safety domain** encompasses all safety-related functionalities as defined in chapter 3. It is continuously available, therefore does not hold capacity for interventions directly on the virtual machine except for reducing the number of active communication channels.

The **quality-managed (QM) static domain** caters to non-safety-critical but essential functionalities. Here, detection of manipulation may prompt the halting of the VM, followed by the initiation of a fallback one. To contain the attack, ongoing service data is not transferred to the fallback VM, which equals a service reset. Moreover, a fallback VM with limited connectivity can be established to prevent an easy replication of the attack while preserving essential QM functions, e.g. related to infotainment systems.

The **QM on-demand domain** facilitates the remaining services that do not require availability, as well as advanced features like edge computing and actual on-demand features. It can be suspended, resulting in the termination of all services. While this inevitably impacts user experience, it serves as a necessary trade-off in case of manipulation.

## 6. Deep dive: the security domain

The security domain plays a critical role within the cross-domain approach, as it is the host of important security components. It is responsible for policies governing access and updates, as well as potentially overseeing the platform itself. The domains are connected via a virtual network in a star topology, as depicted in figure 3, with the security domain at the center. This ensures that data flowing from the QM

devices to the safety domain (and vice versa) must pass through this “virtual gateway”. A wide range of security mechanisms can be applied to this traffic, from a simple firewall over deep package inspection to a complete protocol breaker. All mechanisms report to a central host-based intrusion detection system.




Figure 3: The communication flow between domains (star network) and the technologies used to provide security functions while ensuring performance.

This approach ensures that an adequate level of security is obtained to protect the safety domain. Nevertheless, the required computational overhead can be significant. The safety domain often requires fast response times (e.g. to operate on the CAN Bus) with low latency – typically in the 10 to 20 microsecond range and real-time execution. To guarantee this quality of service for security operations, a hardware safety module (HSM) may be exclusively associated with the safety domain. However, it is logically under the management of the security domain, which is responsible for the distribution and updates of keys stored within it.

While the security domain itself is a regular VM, it will potentially make use of a supporting cast of technologies. HSM Cores are well understood in the automotive industry and provide robust, safety-certified, cryptographic functions. A trusted execution environment (TEE)<sup>3</sup> can provide similar cryptographic features. However, its power lies in its flexibility. The HSM Core is a hardware core within the System-on-Chip (SoC) design. The TEE, in turn, is a combination of a software operating system and hardware security

features of a modern Central Processing Unit (CPU). This enables a TEE to leverage all the features of the CPU – from multiple high-performance cores to large address spaces and peripheral access. In many ways, the TEE can be considered as another VM, albeit with hardware isolation in the form of separate processor states and registers, and isolated memory both in the CPU and across the bus.

We envision the security core using the TEE for many of its functions. For example, policy evaluation related to software updates or inter-domain communication could take place entirely within the TEE to limit access by attackers. Similarly, the TEE can store larger amounts of data, enabling both complex key rotation and access policies, while providing protected storage for audit logs or crash data. This architecture offers significant flexibility and even allows the distribution of security-sensitive tasks between an HSM and the TEE. An application running in the QM static or QM on-demand domains could be granted access to specific application running in the TEE to provide direct access to its own security functions, for example to decode DRM-protected videos.

Finally, the presented segmentation into domains and the security technologies (HSM, TEE, and Hypervisor) allow users to build a Zero Trust platform, which...

... establishes authenticity, confidentiality, and/or integrity protection on a service level, in particular leveraging virtual network channels.

... protects security-critical assets and ensures that access is limited to where it is needed via the HSM.

... protects access and managing rights via TEE and hypervisor.

... monitors the traffic and the status of the entire system via an IDS.

## 7. Use case of a real-life scenario

To illustrate the advantages of the framework, we have chosen an example that particularly emphasizes the origin of the approach in the IT industry. A denial-of-service (DoS) attack is a quite common type of cyber threat that aims at making a device unavailable to its users and unable to maintain its functionality by flooding it with traffic. Whether an IT specialist or the operator of a small website, every person working with networks is familiar with this type of attack, which now also presents a growing challenge in the automotive sector.

As shown in figure 4, we assume that the attackers were able to take over the application in VM3 (e.g. a telematics service), but the actual goal of the attack is the functionality in VM2. In other words, the attackers need to escalate from VM3 to VM2. The special thing here: no further vulnerability is needed to achieve this escalation. The attack relies on the fact that both VM2 and VM3 have legitimate access rights to use the HSM. The functionality in VM2 in this example would rely on the HSM to conduct secure onboard communication (e.g. an immobilizer function or any other signal that affects driving-relevant functionalities). Since VM3 has legitimate access to the HSM, it simply sends so many requests that VM2 cannot perform the secure onboard communication at




Figure 4: Cyberattack via an app on VM3, affecting other VMs.

all or not at the right time, effectively isolating it from the remaining ECUs in the car.

To build our cross-domain security system, we use the setup shown in figure 5. In our example, the QNX hypervisor uses a master virtual machine (VM0). Since it has full control over the guest VMs, VM0 is used for the security domain. It has one CPU core assigned exclusively to it to enable all traffic to be routed and inspected by the security solution. Using dedicated virtual Host2Guest network interfaces, the QM VM (VM3) that holds the telematics service, as well as the safety VM (VM2), which holds the immobilizer service, are connected to the security domain.

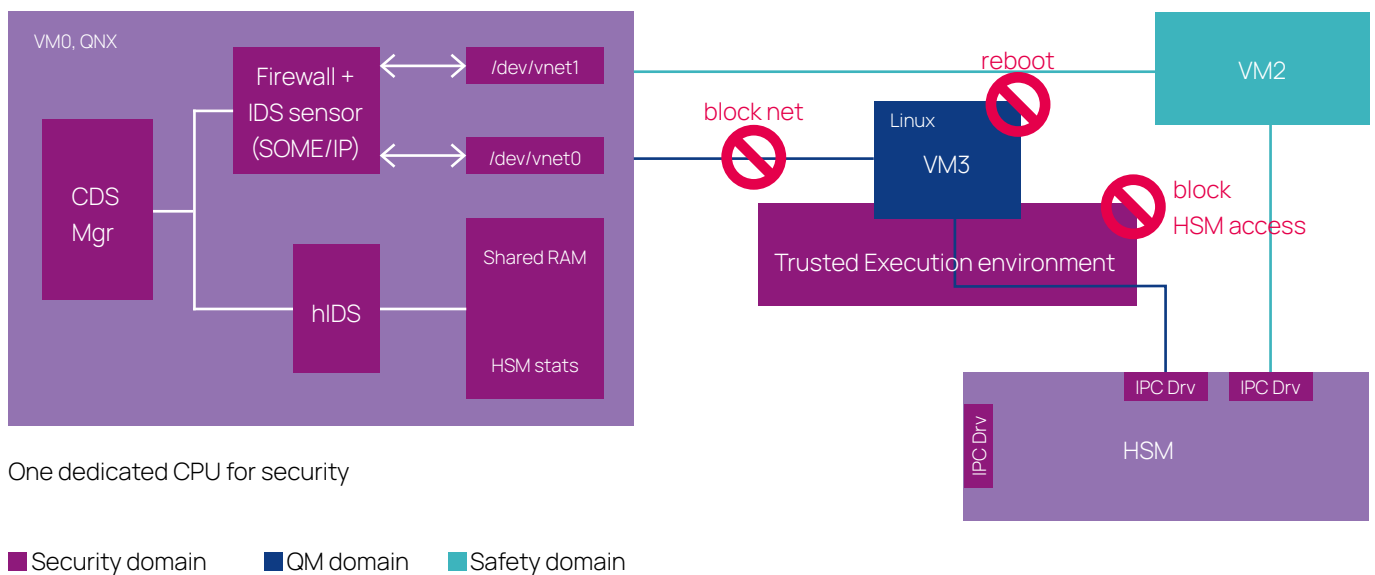


Figure 5: Cyberattack scenario with layered defense responses and bypass over the safety VM.



Both VM2 and VM3 have access to the HSM. By introducing the TEE to connect VM3 to the HSM, all access requests to the HSM can be reported to the central IDS in VM0, enabling an easy detection of the DOS attack. In addition, the TEE is able to enforce access policies, e.g. to limit the access rate or prohibit access entirely. In other words, this setup enables us to react to the attack on the service level. The advantage consists in the high precision combined with a limited impact on the overall functionality. However, it requires understanding and configuring all possible attack scenarios, which cannot be achieved.

To gain a broader reaction capability that does not rely on understanding the attack scenario in detail, we introduce the Cross Domain Security Manager (CDS) which has control over the hypervisor and its network configuration, including the firewall. In case of an attack, the CDS can reboot VM3 and – assuming that the attackers were not able to persist their attack (i.e. bypass secure boot) – effectively get rid of the attackers. To avoid an easy repetition of the attack, the telematics service can be disabled after the VM reboot. This facilitates a reaction to the attack on hypervisor level.

In addition, the CDS can block the network port of the telematic service to the outside world. If the attacked service is known, this setup enables a reaction to the attack on network level.

We assume that the attack was carried out via a random app on VM3. Figure 4 shows how the attack can also affect other vehicle functions located on neighboring VMs, as the flood of requests seeps through to the HSM and paralyzes the system when no proper security measures are in place. Our framework includes detection and layered response actions by the IDS and security domain, which also allow the further functionality of the vehicle computer. In the best case, drivers do not notice the attack at all or only to a limited extent.

As we see in figure 5, the safety VM is unaffected and can keep maintaining vehicle functionalities throughout the attack. The framework with the separate domains therefore not only enables several layers of defense in the event of an attack. It is also designed to maintain as many functionalities as possible, especially those relevant to safety. In contrast to a system with one domain, attacked applications can be decoupled or made functional again via fallback.



## 8. Conclusion

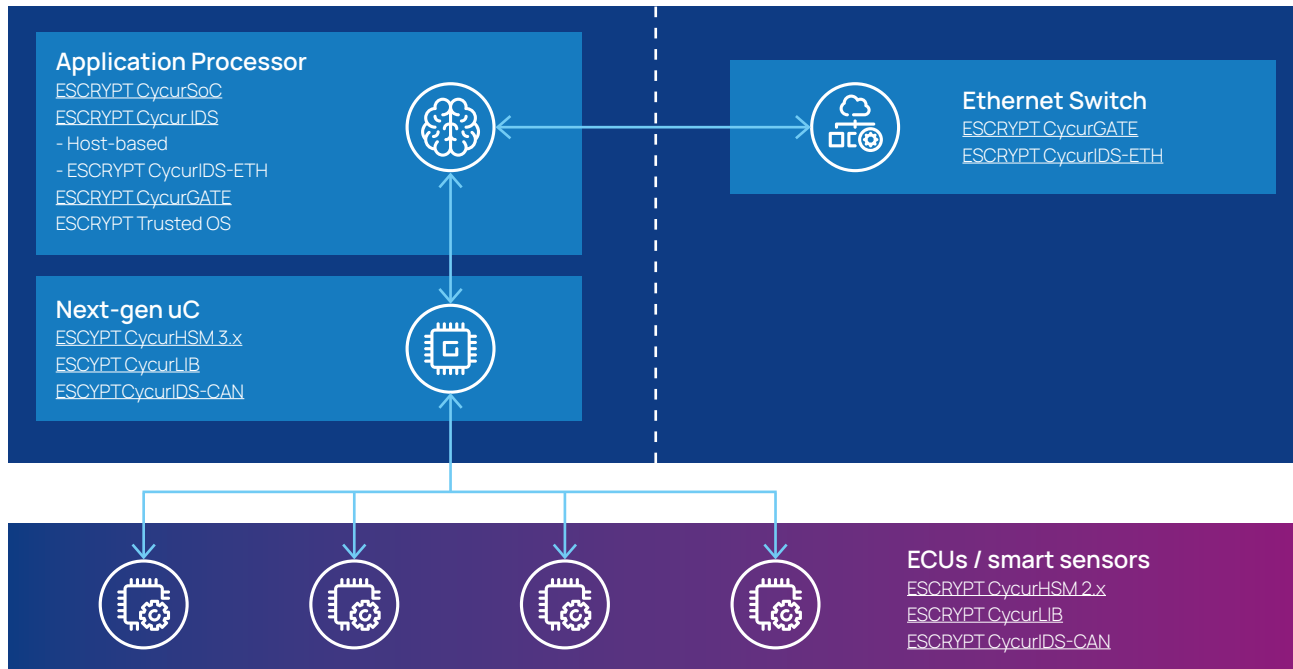
Security is paramount in the rapidly evolving automotive industry. Our proposed “automotive cross-domain security” framework offers a robust approach to securing the vehicle computer, leveraging insights from diverse industries to ensure flexibility without compromising on safety or security. It transfers established best practices of virtualization to split complex systems into isolated units from enterprise/cloud networks to automotive systems. One great benefit of the approach is the possibility to eventually implement a Zero Trust platform, i.e. the next evolutionary step in automotive cybersecurity, within the vehicle computer. This enables manufacturers to set the course for the future by developing an optimally secured vehicle architecture that can withstand any threat – for the maximum safety of all road users.



## ESCRYPT Vehicle Computer Security Suite

With our holistic ESCRYPT Vehicle Computer Security Suite, you can prepare the ground for a cross-domain approach. The solution portfolio comes with plug-and-play packages for specific operating systems and ready-

to-use security sensors. It is customizable by connecting arbitrary sources or native applications with our IDPS module. The fully future-proof portfolio at a glance:



## References

- 1) National Cyber Security Centre: Security principles for cross domain solutions, accessed 2025/02/07, <https://www.ncsc.gov.uk/collection/cross-domain-solutions>
- 2) James E. Smith, Ravi Nair: Virtual machines – versatile platforms for systems and processes. Elsevier, 2005
- 3) Busch, Marcel: On the Security of ARM TrustZone-Based. Dissertation at University Erlangen-Nürnberg, 2020



## About ETAS

Founded in 1994, ETAS GmbH is a wholly owned subsidiary of Robert Bosch GmbH with a local presence in all major automotive markets in Europe, North and South America, and Asia.

ETAS offers comprehensive solutions for the realization of software-defined vehicles in the areas of software development solutions, vehicle operating system, vehicle cloud services, data acquisition and processing solutions, integrated customer solutions and cybersecurity.

As industry pioneers in cybersecurity, we assist our customers in managing cybersecurity-related complexities, reducing cyber risks, and maximizing their business potentials with a proven on- and offboard portfolio of software products and professional security services.

ETAS automotive security solutions are safeguarding millions of vehicle systems around the world – and are setting standards for the cybersecurity of software-defined vehicles.



## Contact information

### **Tobias Klein**

tobias.klein@etas.com

[www.etas.com/we-secure-the-future](http://www.etas.com/we-secure-the-future)

### **Dr. Max Hoffmann**

max.hoffmann@etas.com

[www.etas.com/we-secure-the-future](http://www.etas.com/we-secure-the-future)



All information provided is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and up-to-date information, there can be no guarantee that this information is as accurate as it was on the date it was received or that it will continue to be accurate in the future. No one should act upon this information without appropriate professional advice and without thoroughly examining the facts of the situation in question.

© ETAS GmbH. All rights reserved.

Last updated: 02/2025

### **ETAS GmbH**

Borsigstraße 24, 70469 Stuttgart, Germany  
T +49 711 3423-0, [info@etas.com](mailto:info@etas.com)

Are you interested in  
ETAS products or solutions?  
Please visit **[www.etas.com](http://www.etas.com)**

Or follow us on social media:

